

**US Army Corps
of Engineers.**

Engineer Research and
Development Center

Automated Cartographic Feature Attribution Using Panchromatic and Hyperspectral Imagery

June 2000

Carnegie Mellon University:

David M. McKeown, Jr. , J. Chris McGlone,
Steven Douglas Cochran, Jeffrey A. Shufelt,
Wilson A. Harvey, and Daniel A. Yocum



REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) . June 2000		2. REPORT TYPE Second Annual Technical Contract		3. DATES COVERED (From — To) 1 Apr 1998 — 31 Mar 1999	
4. TITLE AND SUBTITLE Automated Cartographic Feature Attribution Using Panchromatic and Hyperspectral Imagery: DARPA/APGD Yearly Report 1998-1999				5a. CONTRACT NUMBER DACA76-97-K-0004	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHOR(S) David M. McKeown, Jr., J. Chris McGlone, Steven Douglas Cochran, Jefferey A. Shufelt, Wilson A. Harvey, Daniel A. Yocum				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Digital Mapping Laboratory Computer Science Department Carnegie Mellon University 5000 Forbes Avenue Pittsburgh, PA 15213-3891				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Engr. Res. Devl. Center Defense Advanced Research Projects Agency Topographic Engineering Ctr 3701 North Fairfax Dr. 7701 Telegraph Rd. Arlington, VA 22203-1714 Alexandria, VA 22315-3864				10. SPONSOR/MONITOR'S ACRONYM(S) ERDC/TEC CR-00-3	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approval for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report summarizes the primary accomplishments made during the second year of the DARPA Automated Population of Geographic Databases (APGD) program. The first part summarizes work in acquiring and registering the Hyperspectral Digital Imagery Collection Experiment (HYDICE) data. It also describes experiments in classifying the HYDICE data using standard maximum-likelihood techniques, and explores the development and results of the new spectral angle mapper. The second part covers work begun on aggregating the surface material maps to reduce the number of polygons required without degrading classification accuracy.					
15. SUBJECT TERMS block adjustment, linear pushbroom, geometric constraints, HYDICE/hyperspectral imagery, surface material classification, cartographic feature extraction, data fusion					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Lauretta Williams
U	U	U	UU	23	19b. TELEPHONE NUMBER (include area code) (703) 428-6728

Table of Contents

TITLE	PAGE
List of Figures	iv
List of Tables	iv
Preface	v
1. Introduction	1
2. HYDICE Data Acquisition	1
3. HYDICE Block Adjustment	1
3.1 Mathematical model	2
3.2 Experimental plan	2
3.3 Results and evaluation	3
4. Automated Analysis of HYDICE imagery	4
5. Classification Using Spectral Angle Mapping	5
6. Surface Material Map Aggregation	6
7. Data fusion	8
7.1 Feature extraction systems	9
7.1.1 Building Extraction Using PIVOT	9
7.1.2 Stereo Elevation Determination	10
7.2 Improving Building Extraction Performance by Data Fusion	10
7.3 Improving Material Classification Performance by Data Fusion	11
8. Road Network Extraction and Attribution Using HYDICE	14
8.1 Road network extraction	14
8.2 Automated road network attribution	14
9. Conclusions	16
References	17

List of Figures

1	HYDICE test images with tie points, check points, and constrained lines.	2
2	Median absolute XY check point error, meters.	3
3	RADT9 test area from flightline 4.	4
4	Surface material classification in RADT9.	4
5	Surface material classification using the spectral angle mapper in test scene RADT5.	6
6	Surface material classification maps over Fort Hood, TX.	7
7	Aggregation of surface material classification map in CHAFFEE.	8
8	Surface material attribution in a visualization database using an aggregated classification map.	8
9	Aggregation of surface material classification map in RADT5.	9
10	Fort Hood RADT9 test image with extracted edges.	12
11	PIVOT result for RADT9.	12
12	Stereo-based ROIs derived from elevation map.	12
13	Classification-based ROIs derived from classmap.	12
14	Legend for surface material classifications.	12
15	PIVOT results using stereo-based ROI fusion.	13
16	PIVOT results using classification-based ROI fusion.	13
17	Initial projection of PIVOT models to classmap.	13
18	Translation of building polygons using discriminant analysis.	15
19	Surface material assignment after translation.	15
20	Refined classmap using discriminant analysis.	15
21	Final results for surface material assignment.	15
22	Road network extracted using the surface material classification generated from the HY-DICE data.	16

List of Tables

1	Point test cases.	3
2	Fine to coarse class grouping.	4
3	RADT9 top 5 confusion pairs.	5
4	RADT9 coarse classification error matrix.	5

Preface

This research was sponsored by the Defense Advanced Research projects Agency (DARPA) and monitored by the U.S. Army Topographic Engineering Center (TEC), under Contract DACA76-97-K-0004, titled, "Automated Cartographic Feature Attribution Using Panchromatic and Hyperspectral Imagery." The DARPA point of contact is Mr. George Lukes, and the TEC Contracting Officer's Representative is Ms. Laretta Williams.

Automated Cartographic Feature Attribution Using Panchromatic and Hyperspectral Imagery

DARPA/APGD Yearly Report 1998–1999

1. Introduction

This report summarizes the primary accomplishments made during the second year of the Defense Advanced Research Projects Agency (DARPA) Automated Population of Geographic Databases (APGD) program.

Surface material information is of interest to us both for cartographic feature extraction (CFE), to generate feature hypotheses, or to refine features generated by other CFE systems, and for visual simulation to select realistic visual textures. Prior to this contract, late in 1995, we organized a hyperspectral data acquisition using the Naval Research Laboratory's (NRL) Hyperspectral Digital Imagery Collection Experiment sensor (HYDICE) system over Fort Hood, TX. This acquisition resulted in hyperspectral data with a nominal 2-m ground sample distance collected with 210 spectral samples per pixel. These data formed the basis of our program of research under APGD.

Sections 2 and 3 summarize the work in acquiring and registering the HYDICE data. Section 4 describes experiments in classifying the HYDICE data using standard maximum-likelihood techniques, while Section 5, explores the development and results of the new spectral angle mapper.

While we can generate very detailed surface material maps, these are sometimes too detailed for simulation database construction; Section 6 covers the work we have begun on aggregating the surface material maps to reduce the number of polygons required without degrading classification accuracy.

Fusion has been the main goal of this work; in Section 7, we describe fusion experiments involving various combinations of feature detectors and combination modalities.

2. HYDICE Data Acquisition

The collection of data at Fort Hood included both airborne imagery and ground truth measurements. The image acquisition included hyperspectral imagery collected by the HYDICE sensor system and natural color film shot by a KS-87 frame reconnaissance camera. The spectral range of the HYDICE sensor extends from the visible to the short wave infrared (400–2500 nm) regions, divided into 210 channels with nominal 10-nm bandwidths.

Nine HYDICE flightlines, each 640 m wide (cross-track) and 12.6-km long (along-track), were flown over Fort Hood's motor pool, barracks, and main complex areas from an altitude of approximately 4,000 m above ground level. After each flightline, the HYDICE sensor was flown over and imaged a six-step (2, 4, 8, 16, 32 and 64 percent) gray-scale panel, providing in-scene radiometric calibration measurements for each flightline. Prior to the start of the HYDICE flight collection, several ground spectral measurements were made for each gray level panel in an attempt to characterize its mean spectral reflectance curve. A more detailed description of the HYDICE sensor system, Fort Hood image acquisition, and ground truthing activities can be found in [Ford *et al.*, 1997].

3. HYDICE Block Adjustment

Data fusion requires accurate registration between all data types. Our approach to HYDICE registration, based on block adjustment with straight line constraints, has been discussed previously [McGlone, 1998; Ford *et al.*, 1997]. Current experimentation has focused on optimizing the accuracy of the block adjustment and characterizing its results. This section briefly describes experiments in comparing platform models, the effects of varying tie point densities, and the effectiveness of including straight line constraints in the solution.



(a) 4.3.



(b) 5.3.

Figure 1: HYDICE test images with tie points (diamonds), check points (crosses), and constrained lines.

3.1 Mathematical model

The mathematical model has several different parts; the sensor model, which describes the imaging geometry of the linear pushbroom sensor, the platform model, a representation of the aircraft position and orientation with respect to time, and the block adjustment incorporating the geometric (straight line) constraints. In this case the sensor model is based on the collinearity equations, modified to reflect the fact that each line of a pushbroom image is an independent, one-dimensional image [McGlone, 1998].

The platform model describes the behavior of the orientation parameters as a function of time or line number. Two different models were studied in this work, the polynomial model and the interpolative model. In the polynomial platform model, the value of each orientation parameter at a particular line is written as a polynomial function of line number. The interpolative model, on the other hand, stores the orientation parameters of reference lines at regular intervals, then calculates the parameters of intervening image lines by polynomial interpolation.

The bundle block adjustment is performed using an object-oriented photogrammetry package [McGlone, 1995] which allows the use of images with different geometries and the rigorous incorporation of straight line constraints. Straight lines in the scene are measured in each image, in order to provide additional strength to the solution.

3.2 Experimental plan

For the experiments described in this report, a small sub-block of the available data is being used. This includes two sidelapping 1280-line HYDICE images, four KS-87 images (1.0 meter GSD), and four RADIUS vertical mapping images (0.3 meter GSD). Tie points between the HYDICE images and the frame images were established by manual measurement, with all tie points being measured on at least two frame images. Straight lines also were measured manually on at least two frame images. The two HYDICE images used are shown in Figure 1. Tie points for the heavy density case (described below) are shown as diamonds while check points are shown as crosses. The straight lines used in the solution also are shown.

Three levels of tie point density were established, as shown in Table 1. The same 37 check points were used for each experiment. Check points that appear on both HYDICE images are counted twice, since they are treated independently. All measured object-space straight lines were horizontal and were constrained to be horizontal.

Table 1: Point test cases.

Case	Pts on 4.3	Pts on 5.3	Pts on both
1 (heavy)	18	17	8
2 (medium)	14	12	5
3 (sparse)	6	6	3

3.3 Results and evaluation

Evaluation was done by comparing the calculated world X,Y coordinates of the check points against the values using the frame images. No evaluation was done on the Z coordinate, since the HYDICE sensor has a very narrow field of view (9 degrees) and therefore elevation recovery is weak. For this reason, the Z coordinates of the check points were held fixed in the solution, and points that appeared on both HYDICE images were evaluated as two separate points. In order to gain a better understanding of the characteristics of the solution, three different statistics were calculated: the median absolute deviation, the root-mean-square (RMS) deviation, and the maximum absolute deviation. Since the RMS statistic is extremely sensitive to large outliers, we rely mostly on the median statistics in analyzing the results.

The results of the evaluation runs are given in Figure 2. The interpolative model solution for the sparse point case (3) with no lines and 32-line spacing did not converge, due to weak geometry with the reduced number of points, so no results are given.

Polynomial vs interpolative platform models:

For this data set, the polynomial model generally performed better than the interpolative model without lines, but not as well as the interpolative model with straight line constraints. The interpolative model without straight line constraints degrades more rapidly than the polynomial model as the amount of control is decreased (going from the heavy (1) to the sparse (3) point densities).

Effectiveness of straight line constraints: The inclusion of straight line constraints in the interpolative model solutions improved the results in every case. While decreasing the number of tie points still increased check point error, the results from the runs with sparse points (case 3) are still better than the results for the heavy point density (case 1) without lines. This indicates that straight line constraints can be used both to improve a solution or as an effective substitute for additional tie points; however, adding the straight line constraints to the polynomial model solution made only negligible differences. It may be that the polynomial model, with its more limited flexibility, is unable to use the additional information from the line constraints.

Reference line spacing: Decreasing the reference line spacing for the interpolative model will make the model more flexible by increasing its degrees of freedom. Given enough information to determine the model, it should recreate the platform motion more accurately and give better results. In this case, however, decreasing the reference line spacing generally degraded the results. The additional degrees of freedom were not adequately determined by the available information, and, in fact, the solution using the sparse point density without lines did not converge.

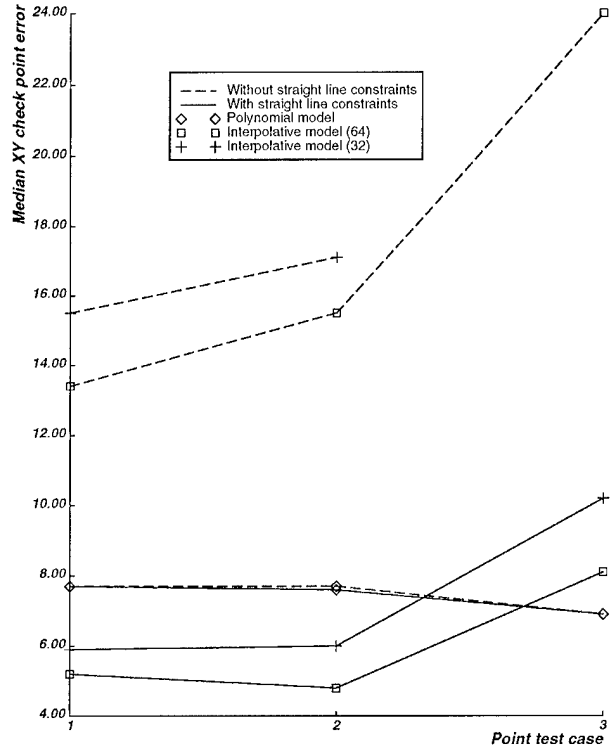


Figure 2: Median absolute XY check point error, meters.

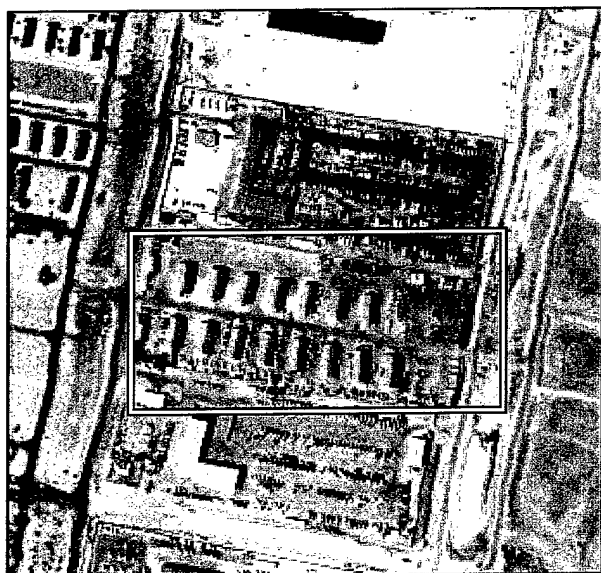


Figure 3: RADT9 test area from flightline 4.

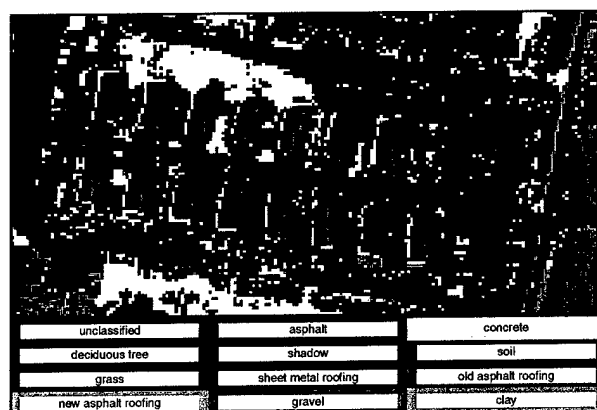


Figure 4: Surface material classification in RADT9.

4. Automated Analysis of HYDICE imagery

Table 2: Fine to coarse class grouping.

Coarse Surface Material	Fine Surface Material
man-made surface	asphalt concrete
bare earth	soil clay gravel
vegetation	grass deciduous tree coniferous tree
water	deep water turbid water
man-made roofing	new asphalt roofing old asphalt roofing sheet metal roofing
shadow	shadow

Due to the volume of image data collected by the HYDICE hyperspectral sensor, these classification experiments used a reduced image dataset. To build on our previous experience with Daedalus (Airborne Thematic Mapper (ATM) imagery, we simulated Daedalus ATM imagery by averaging the HYDICE imagery bands contained within the solar reflective bandpasses of the Daedalus ATM scanner.

Figure 3 shows one of the test areas used in the surface material classification experiments. Manually-selected training sets for the materials listed in the "Fine Surface Material" column of Table 2 were compiled from an earlier section of Flightline 4. A Gaussian Maximum Likelihood (GML) classification was performed using the 10 simulated Daedalus ATM bands and selected training sets. Figure 4 shows a surface material subsection map from the resulting classification, corresponding to the outlined region shown in Figure 3.

The resulting surface material map was evaluated against manually-generated surface material reference data. Overall classification accuracy was 57.9 percent for RADT9. From Table 3, almost 20 percent of RADT9's classification error is associated with confusion among concrete, asphalt, soil, and gravel. Looking at Figure 4, there is breakup of the parking lot into asphalt and concrete sections probably influenced by surface weathering and vehicular traffic. Also, the barrack roofs fluctuate in surface material classification due to illumination changes influenced by the structure of the building roofs.

We also are interested in coarse surface material classification, whereby the fine surface material classes are grouped into more general categories as listed in Table 2. This type of broad categorization is useful in identifying areas containing man-made or natural surface features. Table 4 displays the error matrices for the coarse classification for the test area, with a classification accuracy of 75.0 percent. The majority of the error (10.4 percent) involves man-made surface and bare earth confusions.

5. Classification Using Spectral Angle Mapping

The spectral angle mapper (SAM) generates surface material classification maps by determining the spectral similarity between test and reference spectra. The reference end member spectra are extracted from the hyperspectral imagery and represent the spectral signature of each canonical 'class'. The mean spectral curve of an asphalt parking lot, for example, may be used as the reference end member spectra for the asphalt class. The similarity is determined by measuring the angle between each test spectra and the reference spectra in n-dimensional space, where n is the number of bands available in the imagery. Each test area is then assigned to the reference end member class to which it is most similar, i.e., to the reference class that has the smallest angle with the test spectra. Unlike the maximum likelihood classification discussed earlier, the spectral angle mapper uses the full spectral range (210 bands) of the HYDICE imagery, allowing more information to be used in the discrimination of surface materials.

To evaluate the SAM, surface material classifications were generated with the same set of classes used in the maximum likelihood classifier discussed earlier, based on reference end member spectra created from the same training sets. The resulting surface material maps were evaluated against manually-generated surface material reference data for the two test areas. Classification accuracies for CHAFFEE were 55.9 percent while those for RADT5 were 55.1 percent. Visual inspection of the SAM classification results shown in Figure 5 reveal that classification of vegetation features appears to be quite good. Unlike the results seen in the surface material maps generated by the maximum likelihood classifier, building roof tops have relatively homogeneous classification results. One of the primary sources of confusion in both test areas is between the asphalt and gravel classes, most notably along roads. This confusion is most likely due to mixed gravel and asphalt pixels along the road shoulders.

In order to more efficiently examine the results of surface material classification maps generated with the spectral angle mapper, the Digital Mapping Laboratory has created an interactive SAM tool. The program, IDI_SPANGLE, allows the user to manually delineate a series of reference end members regions within an image, or to specify a pre-selected set of reference spectral information. Surface material classification can then be generated on an area of the image delineated by the user. IDI_SPANGLE also enables the user to observe the spectral signature of reference end members, which can be useful in developing a reference set with adequate class separability. After the SAM classification is completed, the resultant surface material map is displayed. After examining the classification results generated from the selected reference spectra set, the spectral angle calculation can be rerun with a different set of reference end member regions. The ability to rapidly adjust reference end-member spectra and observe the resultant surface material classification maps is a useful tool for studying the effects of varying reference spectra on spectral angle classification.

Table 3: RADT9 top 5 confusion pairs.

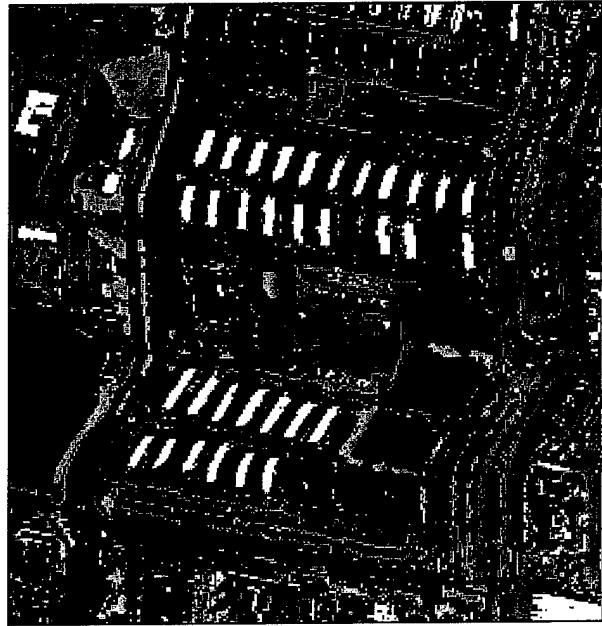
Ground Truth Class	Classification Class	Number Confused	Error
concrete	asphalt	7074	10.3%
soil	gravel	2756	4.0%
grass	soil	2233	3.3%
asphalt	soil	1703	2.5%
concrete	soil	1558	2.3%
Total		15324	22.4%

Table 4: RADT9 coarse classification error matrix.

TEST	REFERENCE						Row Total	Commission Error
	man-made surface	bare earth	vegetation	water	man-made roofing	shadow		
man-made surface	26677	1973	1509	0	1477	608	32244	17.3
bare earth	5129	7299	2444	0	254	58	15184	51.9
vegetation	377	822	14666	0	117	149	16131	9.1
water	0	0	0	0	0	0	0	*
man-made roofing	671	182	285	0	1598	367	3103	48.5
shadow	130	33	425	0	79	1153	1820	36.6
Column Total	32984	10309	19329	0	3525	2335	68482	
Omission Error	17.3	51.9	9.1	*	48.5	36.6		Percent
Overall Accuracy = 51393 / 68482 = 75.0%								



(a) HYDICE short wave infrared representation of test scene RADT5.



(b) Surface material classification map for test scene RADT5.

Figure 5: Surface material classification using the spectral angle mapper in test scene RADT5.

6. Surface Material Map Aggregation

Current research in the Digital Mapping Laboratory focuses on providing surface material attribution for visualization databases. Problems arise, however, when attempting to use surface material classification maps generated from HYDICE imagery. Figure 6 shows three surface material classification maps available for use in visualization databases. The Interim Terrain Data (ITD), Figure 6a, provides broad areal coverage but has coarse spatial resolution. Surface material classification maps generated from HYDICE data with a 2-meter GSD, Figure 6b, show much more spatial detail. This detail, however, is often too high to feasibly use in a visualization database. Aggregation of the HYDICE derived surface material maps, Figure 6c, could provide more detailed surface material maps than ITD derived maps, without the high cost of non-aggregated HYDICE material maps.

In order to more conveniently work with the raster format surface material maps, they were converted to a polygonal format. Aggregation of the surface material classification maps was then conducted by first removing all 'tree' regions from the classification map. A separate file was created with locations of these tree regions for later use in visualization database generation. Following removal of the tree polygons, all classification regions within the image were examined to see if they fell below a specified area threshold. If the regions were small enough, they were merged into surrounding polygons. The selected regions were then merged with the neighboring region that had the greatest shared perimeter with that region.

Figure 7 shows the result of several aggregation tests on the CHAFFEE test scene. The classification map used in these studies was generated with a maximum likelihood classifier. The experiment was completed using minimum area thresholds of both 120 m², Figure 7b, and 800 m², Figure 7c. Prior to aggregation, the CHAFFEE surface material classification map contained 5,431 polygonal regions. Aggregation of all regions below 120 m² reduced the polygon count to 200, while aggregation of all regions below 800 m² trimmed the polygon count to 33. A reduction of the polygon count within the surface material classification map, while maintaining general classification accuracy, allows more efficient surface material attribution during the generation of visualization databases.

An example of surface material attribution of a visualization database with an aggregated surface material map can be seen in Figure 8. The tree population in this example was generated by using point

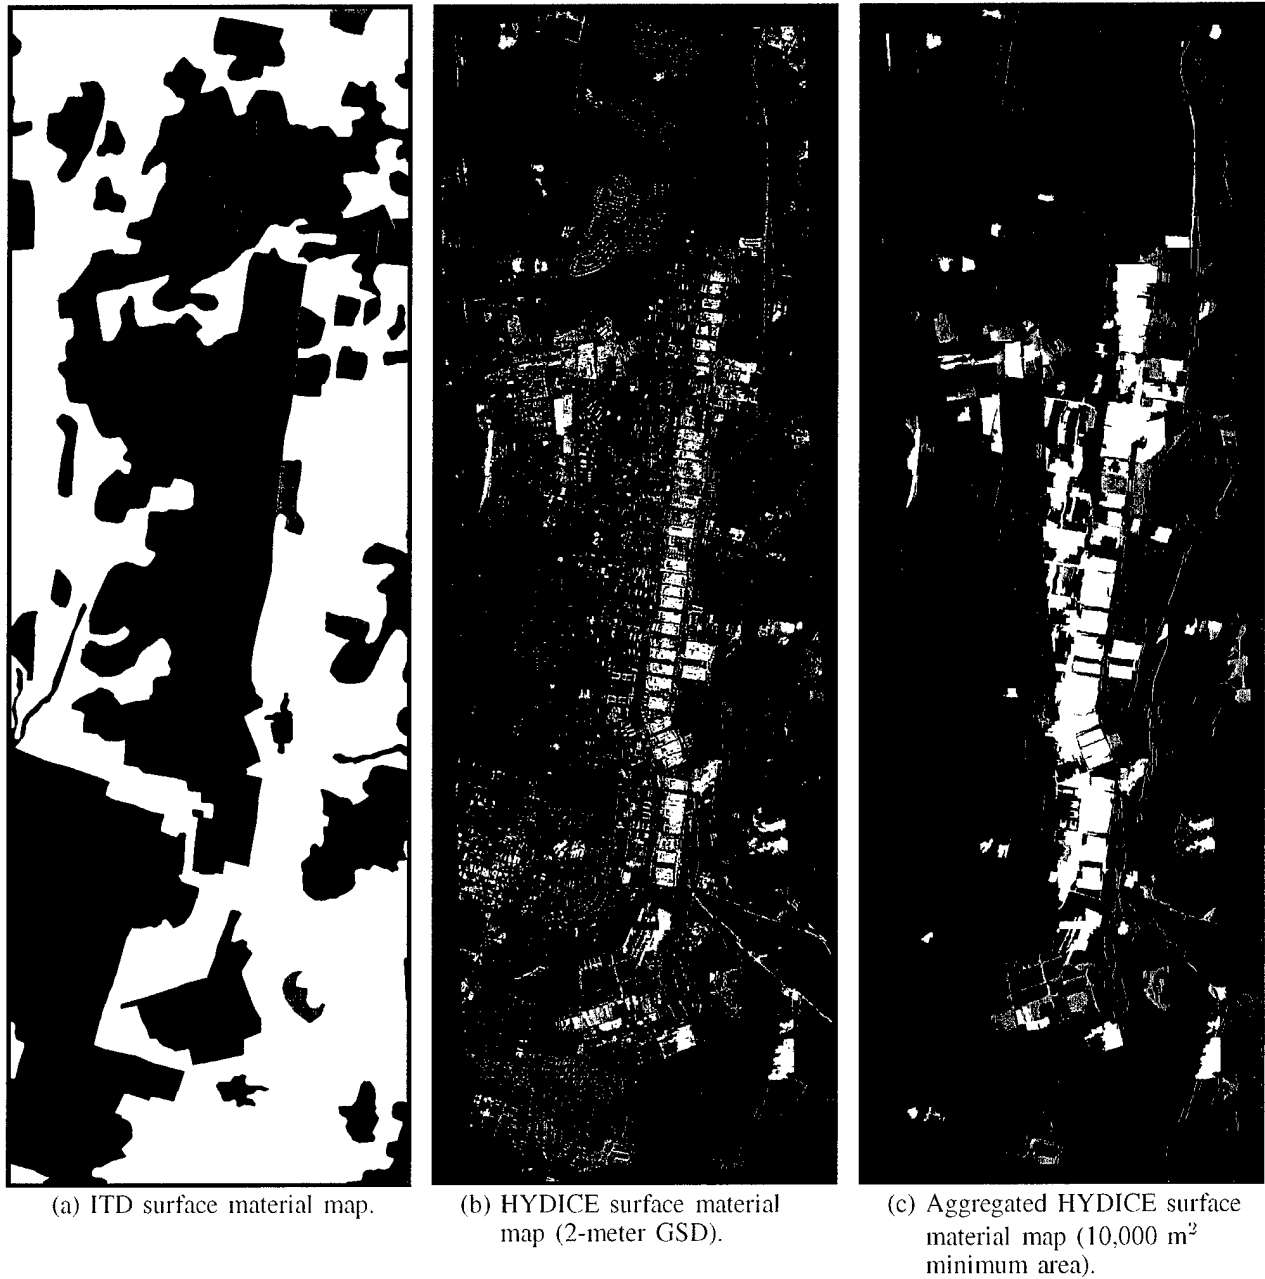
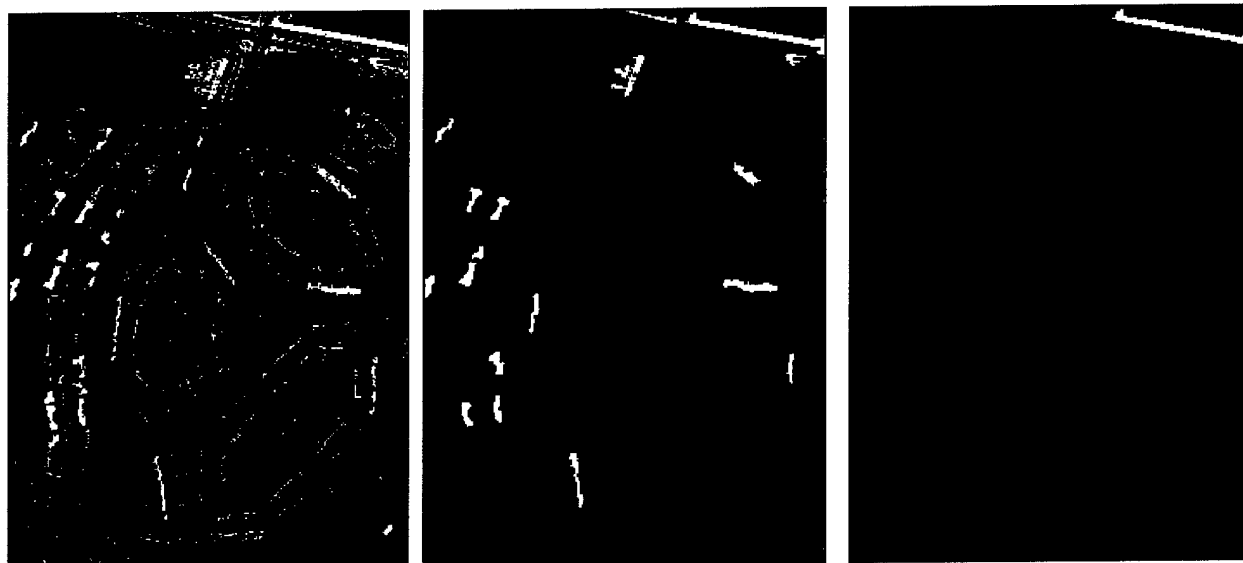


Figure 6: Surface material classification maps over Fort Hood, TX.

locations from the trees removed during aggregation. As the figure shows, an aggregated surface material classification map can be used for attribution of the surface material types in a visualization database without requiring an infeasible number of polygons to represent the material attributes.

Aggregation of surface material classification maps also can be used to generate more homogeneous classification maps for use in information fusion with other cartographic feature extraction systems. Figure 9 compares the original maximum likelihood classification results against both an aggregated surface material classification map and manually generated reference data. Visual inspection of the results indicates that aggregation removes many of the small, mixed pixel regions from classification results. Percentage accuracy of the classification results also improves. Original classification results have a 65.1 percent accuracy with respect to the reference data, while the aggregation classification map has an accuracy of 67.5 percent.



(a) Original maximum likelihood classification results.

(b) Surface material map with 120 m² aggregation.

(c) Surface material map with 800 m² aggregation.

Figure 7: Aggregation of surface material classification map in CHAFFEE.



Figure 8: Surface material attribution in a visualization database using an aggregated classification map.

7. Data fusion

The common theme throughout our research has been the belief that no single computer vision technique can reliably provide a complete scene reconstruction; thus, to achieve good performance, we need to gather a variety of information, extracted by various processes from multiple images of the area of interest, and synthesize this disparate information into a consistent model. This is the *cooperative methods* approach to cartographic feature extraction.

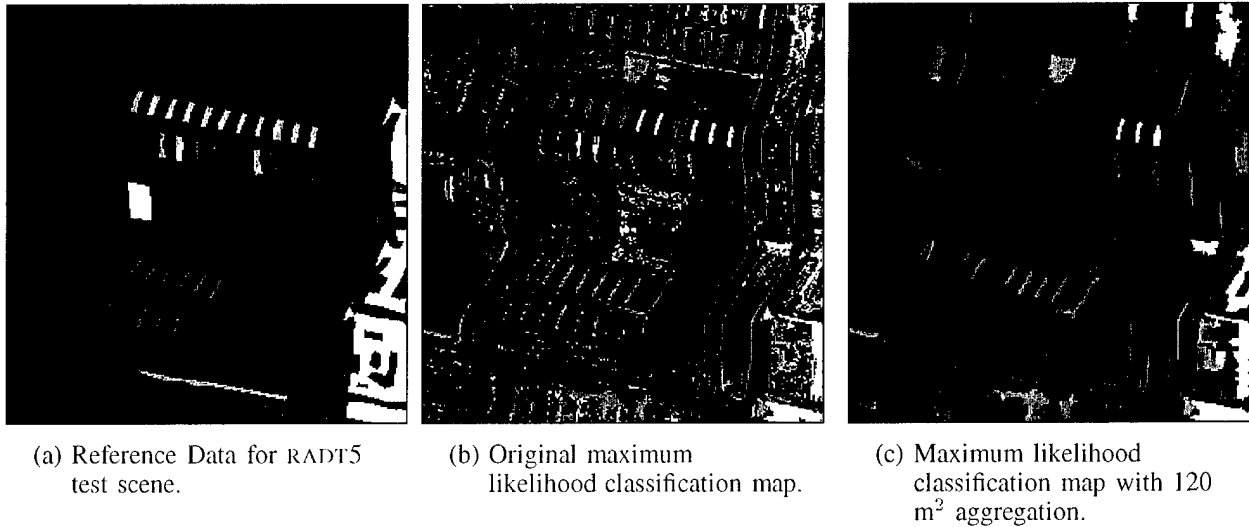


Figure 9: Aggregation of surface material classification map in RADT5.

This leads to the central question: How can we intelligently combine and integrate the different sources of partial information, generated by our feature extraction systems, to facilitate 3-D scene analysis? We seek to improve overall performance both in terms of better quality and faster processing.

There are alternative ways for organizing the 3-D scene reconstruction threads into a combined processing approach. The most basic division is either into a bottom-up (data directed) approach, where the results from the different methods are merged together; or, a top-down (knowledge-directed) approach, where the partial or full results from one source are used to guide or select the processing of other approaches.

We have used both approaches in our research and sometimes combine them in order to maximize the use of the information available from the different systems. This section discusses the individual feature extraction systems and gives examples of the application of fusion techniques to building extraction, surface material classification, and road network extraction.

7.1 Feature extraction systems

In this section, we briefly describe the cartographic feature extraction systems that serve as the basis for our experiments in data fusion. We extract four kinds of features for data fusion:

- Surface material maps obtained from the classification of hyperspectral imagery (discussed in Section 4),
- Digital elevation models derived from stereo panchromatic imagery,
- 3-D building hypotheses generated from single panchromatic images, and
- Road network hypotheses.

7.1.1 Building Extraction Using PIVOT

Perspective Interpretation of Vanishing points for Objects in Three dimensions (PIVOT) is a data-driven fully automated monocular building extraction system developed at the Digital Mapping Laboratory [Shufelt, 1996]. PIVOT is based on two key ideas: first, that photogrammetric knowledge can be exploited at all phases of the building extraction process to improve performance; and second, that buildings can be well modeled by composition of a small set of volumetric *primitives*.

The inputs to PIVOT consist of a panchromatic aerial image, the interior and exterior orientation, and the date and time of image acquisition. From these inputs, PIVOT produces 3-D wireframe representations of the buildings in the image, referenced to geodetic coordinates. PIVOT makes use of the photogrammetric

camera model to detect vanishing points for the image that correspond to the shapes of the primitive volumes to be extracted.

One consequence of this vanishing point-based approach is that the performance of PIVOT depends heavily on the quality of the underlying edge data. This dependence on edge data for feature extraction is not unique to PIVOT; many current building extraction systems rely on clean edge data to extract structure.

7.1.2 Stereo Elevation Determination

Idl.DPCP is a stereo system that can operate either fully- or semi-automatically and generates an object-space elevation estimate from two or more images of a scene. The Digital Photogrammetry Compilation Package (DPCP) stereo matcher was developed at the U.S. Army Topographic Engineering Center [Norvelle, 1981; Norvelle, 1992] and the interface and application of the stereo matcher was developed at the Digital Mapping Laboratory [McKeown *et al.*, 1997].

The input to Idl.DPCP is two or more panchromatic aerial images covering a common area along with the interior and exterior orientation data for each image. From these inputs, Idl.DPCP produces an elevation map in object-space for the scene. Idl.DPCP requires no parameter adjustment or threshold tuning, nor does it require any additional image adjustment. All of these are performed automatically by the system.

If multiple stereo pairs are available, then multiple individual image pyramids are built from coarse to fine where, at each level, the current best elevation estimate is shared amongst all of the stereo pairs for use in the next level of the pyramid match. This provides a significant improvement in accuracy in the final result by removing most of the stereo “blunders” and by decreasing the effects of biasing that might be present on a single stereo pair.

7.2 Improving Building Extraction Performance by Data Fusion

As noted in Section 7.1.1, many building extraction systems are data-driven, beginning with an edge extraction process to supply low-level geometry for the inference of building structure. One of the key difficulties such systems face is focusing their processing on edges and lines that correspond to actual building structure. In this section, we explore two similar approaches to use multisource data fusion to focus edge processing.

Both approaches share the same underlying technique. The line segments produced by edge extraction on a single panchromatic image are filtered through “regions of interest” (ROIs), areas in the image that are believed to have a high probability of containing building structure. Line segments which do not come into contact with any ROI are discarded, and no further analysis is performed on these segments. The remaining line segments are then processed by PIVOT to extract building structure.

The first approach, *stereo-based ROI*, uses a dense elevation map generated from multiple panchromatic images as the data for determining ROIs (Section 7.1.2). These are generated by projecting both the high-resolution stereo elevation results and the low-resolution DEM from object-space to the desired image-space coordinates. Once both sets of elevation information have been projected, they are differenced (to remove low spatial frequency ground elevation) and the contiguous areas with an average differential height of more than 2-meter are designated building hypotheses (differencing after the projection takes care of occluding edges generated by the building). The resultant set of areas includes buildings, trees, and very large vehicles, but serves to dramatically reduce the search space for buildings, especially through walkways and across parking lots. These initial hypotheses are then morphologically expanded by five pixels to allow for stereo “blunders” and for shrinkage due to the elevation threshold; regions of less than 600-m² are excluded.

The second approach, *classification-based ROI*, uses a material classification map generated from multispectral imagery as the basis for generating ROIs. In this approach, ROIs consist of those collections of pixels with classifications corresponding to building roof materials. In addition, since PIVOT analyzes shadow geometry to infer 3-D structure, pixels assigned to the shadow class also are included as ROIs to allow edges formed by shadows to pass through to PIVOT.

In our experiments we have used *weak edge filtering* with ROIs which allows an edge to pass if it comes into contact with an ROI. An alternate method, *strict edge filtering*, only allows an edge to pass if it is contained entirely within an ROI. This method is not used because it is less robust than weak filtering; if ROIs do not completely cover a building, then edges that fall in the gaps of an ROI will be discarded under strict edge filtering, even though they may delineate the underlying building structure. Of course, weak edge filtering can allow extraneous edges to pass, but this is a better alternative for data-driven bottom-up analysis systems such as PIVOT, which require some minimal amount of boundary information to infer structure, and typically employs a hypothesis verification step to reject spurious building models.

Figure 10 shows an aerial image (RADT9) of a set of barracks in Fort Hood, with the results of edge extraction superimposed on the image. Figure 11 shows the resulting building hypotheses generated by PIVOT. While some buildings are properly delineated, many false positives are generated by edges that have geometric structure consistent with buildings, even though these edges lie along road and parking lot boundaries.

Applying the stereo-based ROI method to an elevation map for the Fort Hood scene produces the ROIs shown in Figure 12, with the filtered edges superimposed on the ROIs. Comparing Figure 12 with Figure 10, note that a significant number of edges have been discarded because they do not come into contact with the stereo-based ROIs; in other words, they do not correspond to height discontinuities in the scene. Of course, other tall features in the scene also can generate ROIs and filter edges; many of the edges in the lower right corner of Figure 12 are caused by large trucks in the parking lot.

Figure 13 shows the classification-based ROIs resulting from the classification of test area RADT9 (Figure 4), again with the filtered edges superimposed on the ROIs. Figure 14 shows the color coding used for the surface materials in the classmap. Again, comparing Figure 13 with Figure 10, many edges have been discarded because they do not touch regions with a roof material or shadow class. Edge filtering with these ROIs is not as effective as in the stereo-based ROI case, since there are many isolated pixels with a roof material class that comes into contact with edges.

Figures 15 and 16 show the final PIVOT results using stereo-based and classification-based ROIs, respectively. These compare favorably with the original PIVOT result in Figure 11, in terms of significantly fewer false positive building hypotheses, and in some cases, more accurate building delineations; these results illustrate the strengths of the ROI approach to fusion. To accurately assess performance improvements, however, it is critical to employ a quantitative comparative analysis of PIVOT with and without ROI fusion.

7.3 Improving Material Classification Performance by Data Fusion

In the previous section, information derived from multispectral classification served as a mechanism for improving the performance of building extraction systems for monocular panchromatic imagery. It is natural to ask whether the converse is true: can models extracted automatically from panchromatic imagery be used to refine and improve a classmap derived from multispectral imagery?

One simple approach for using the building polygons generated by PIVOT projects them into the HYDICE image space, selects a representative class for each building polygon, and replaces the class values inside each polygon with its representative class. The basic idea is that the building polygons define regions that have homogeneous surface material composition, and any misclassified pixels within these regions can be corrected by assigning all pixels in the region to an appropriate class. More generally speaking, geometry derived from panchromatic imagery is used to refine classification in multispectral imagery.

Such an approach must handle two kinds of issues: registration issues, which arise due to the limits of projection accuracy of the camera models across different sensor platforms, and classification issues, involving the choice of methods for selecting an appropriate representative material for each building polygon. The simplest approach for handling these issues is to project building polygons through the sensor models directly to the HYDICE imagery without any corrections, and assign each projected polygon the most frequently occurring class inside its perimeter.

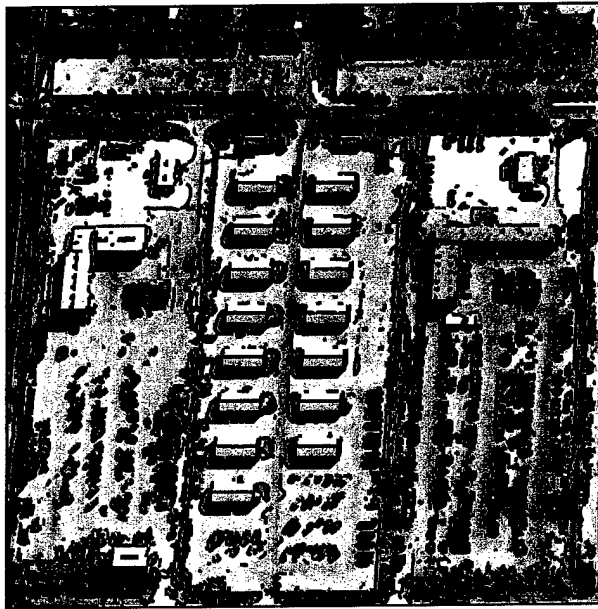


Figure 10: Fort Hood RADT9 test image with extracted edges.

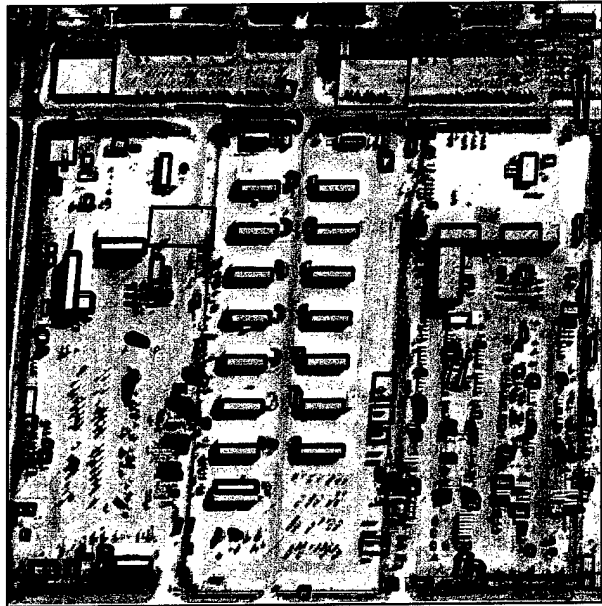


Figure 11: PIVOT result for RADT9.

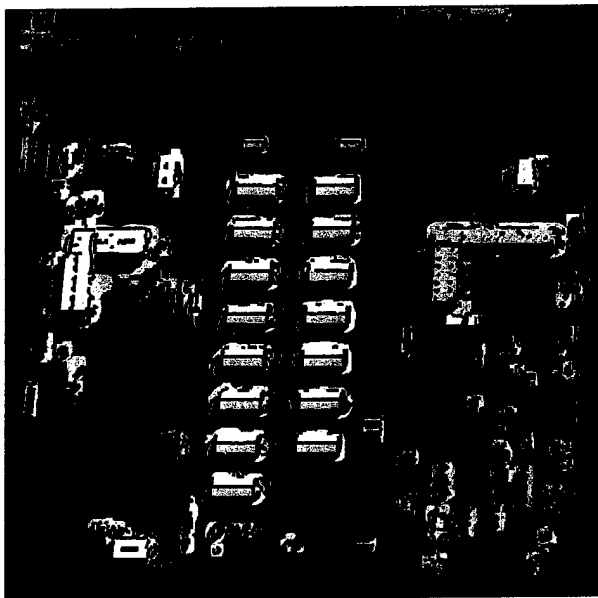


Figure 12: Stereo-based ROIs derived from elevation map.

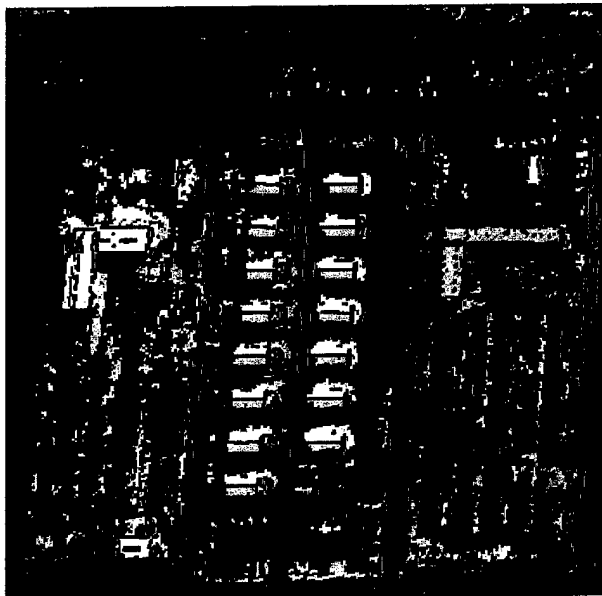


Figure 13: Classification-based ROIs derived from classmap.

asphalt	clay	concrete	coniferous tree	deciduous tree
deep water	grass	gravel	new asphalt roofing	old asphalt roofing
shadow	sheet metal roofing	soil	turbid water	unclassified

Figure 14: Legend for surface material classifications.

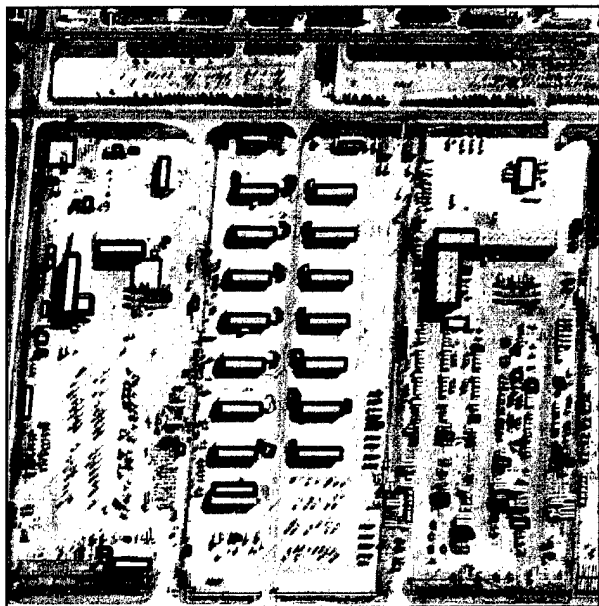


Figure 15: PIVOT results using stereo-based ROI fusion.

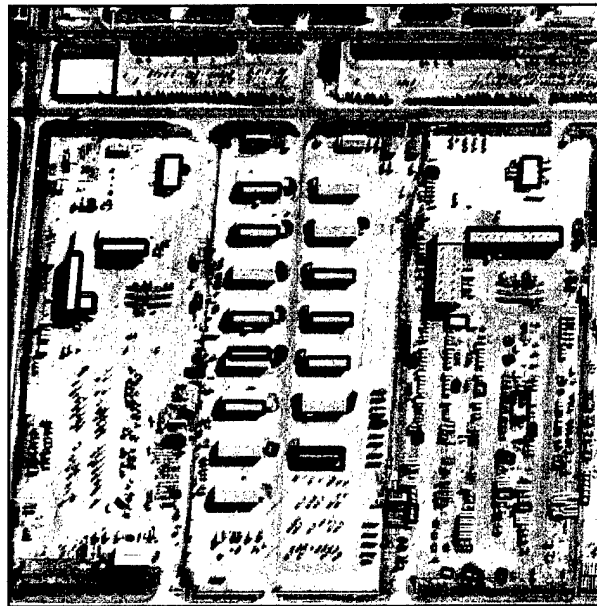


Figure 16: PIVOT results using classification-based ROI fusion.

In the following figures, we use PIVOT with stereo-based ROI to provide building models. Figure 17 shows the result of projecting the models directly to the HYDICE coordinate system; the barracks buildings in the center of the image do not line up with the underlying classification results due to projection error in the sensor model. Some building roofs even overlap shadow and grass classes.

To address these issues, we present new methods for handling the registration and classification problems. Rather than use the final classmap directly, we instead use the *discriminant* values calculated from a GML classifier, which can be treated as weighted distances from a pixel to a class representative in spectral feature space.

The first step is to address the registration problem. After a building polygon and its shadow are projected into the HYDICE imagery, these are shifted vertically and horizontally within a larger window, bounded by the maximum expected projection error, to find the best position. This is determined by a least-squares solution that minimizes the discriminant distance between the roof polygons and the roof material class representatives, and the discriminant distance between the shadow polygon and the shadow class representative. This approach models the registration error as a translation in HYDICE image space.

Figure 18 shows the results of applying the new registration method. In comparison with Figure 17, the barracks buildings now generally line up well with the underlying regions in the classmap. One exception is the L-shaped building at the top center of the image, for which one wing of the building has been incorrectly translated to the right. This approach still depends, albeit less directly, on the quality of the class representatives, since these determine the discriminant distances used for the least-squares fit. Figure 19 shows the results of choosing the maximally occurring class for each of these shifted

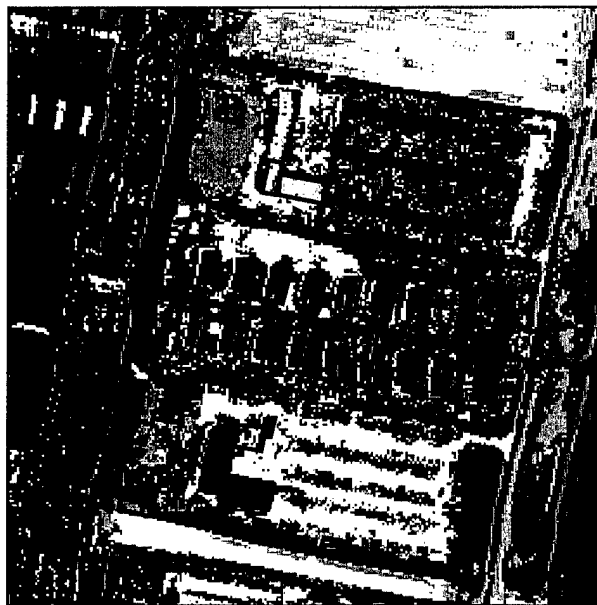


Figure 17: Initial projection of PIVOT models to classmap.

polygons. Note that the class attributions for the building models are now significantly better, due to the improved registration; some buildings, however, still have incorrect roof material assignments.

To improve the material assignment, we again make use of the discriminant values. Rather than choose the maximally occurring class in a polygon, we instead select the class that has the minimum mean discriminant distance from the pixels in the polygon, limiting the selection to roof material classes. These conditions ensure that we not only select a legal roof class, but that we also select the roof class with the best overall fit to the region defined by the polygon.

Figure 20 shows the same projection result as in Figure 18; the difference is in the coloring inside each polygon, which is now determined by the minimum mean class discriminant for each polygon. Figure 21 shows the final result: compared with Figure 19, several buildings have had their roof materials modified.

In this last example, two kinds of fusion have taken place. In addition to automatically assigning surface material attribution to the roof polygons of building models (Figure 21), the classmap itself has been refined by assigning the interior pixels of each building polygon with a homogeneous class, determined by discriminant analysis.

8. Road Network Extraction and Attribution Using HYDICE

Hyperspectral image data and derived surface material maps can provide powerful cues to road network extraction systems. This section presents the results of using feature extraction techniques on HYDICE image data and derived surface material maps and also a novel use of surface material maps to refine USGS Digital Line Graph (DLG) map data.

8.1 Road network extraction

In addition to attribution, surface material information can be used to provide clues for detection of road features. This information can be especially helpful in cluttered areas, such as suburban housing areas where the roads may be obscured by vegetation and/or shadows. We present some initial experimental results applying our road network extraction systems to HYDICE and HYDICE-derived image data.

Given a surface material classification of the scene, we segment the scene based on the classification results, then extract those regions that have surface material types relevant to the road features that we want to delineate, *e.g.*, asphalt, gravel, and concrete. After simplification and smoothing, these regions can be used to generate segments by doing thinning and connected component extraction. Widths can be assigned by overlaying each segment on the surface material map and, for each point, calculating the width of the corresponding region in the surface material map. We then extend these segments with our composable road tracker [McKeown *et al.*, 1998] using the original HYDICE image data as input to the tracker. Finally, we generalize the tracked roads, bridge gaps, then turn the vectors into a road network.

8.2 Automated road network attribution

Although the specification supports it, few Digital Line Graph (DLG) transportation data sets have width or surface material attributions. Using surface material maps generated from HYDICE classifications, we can augment an existing DLG road network with detailed width and surface material information.

We begin by projecting and clipping the DLG data into the desired image space. We then guess a width for each road segment. This can be done very approximately by assigning all the roads the same "average" width, or by running an automated road finder on the surface material map and matching the various road seeds to portions of the DLG road network. Next, we deal with any registration errors by matching the road network to the surface material map by using a discriminant optimization technique [Ford *et al.*, 1999]. In order to do this accurately for roads, we preprocess the network by creating polygons around each intersection. This step is necessary because roads tend to be symmetric objects; the intersection polygons will tend to be asymmetric and, therefore, match more uniquely. This registration refinement step yields a set of translations for each intersection polygon that then must be propagated back to other points

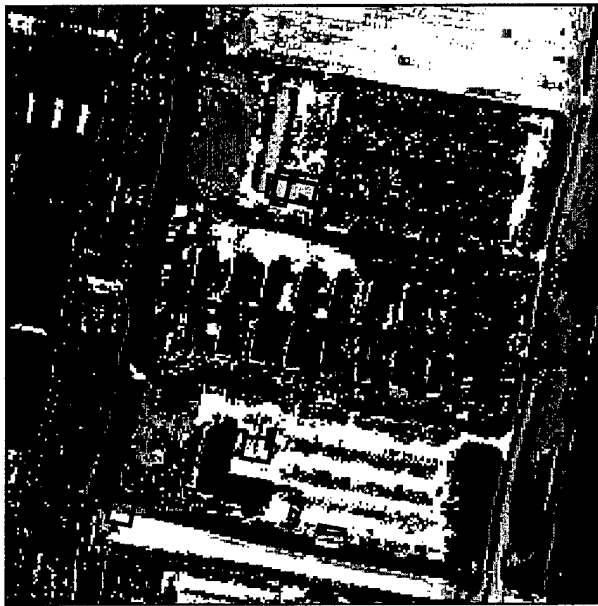


Figure 18: Translation of building polygons using discriminant analysis.

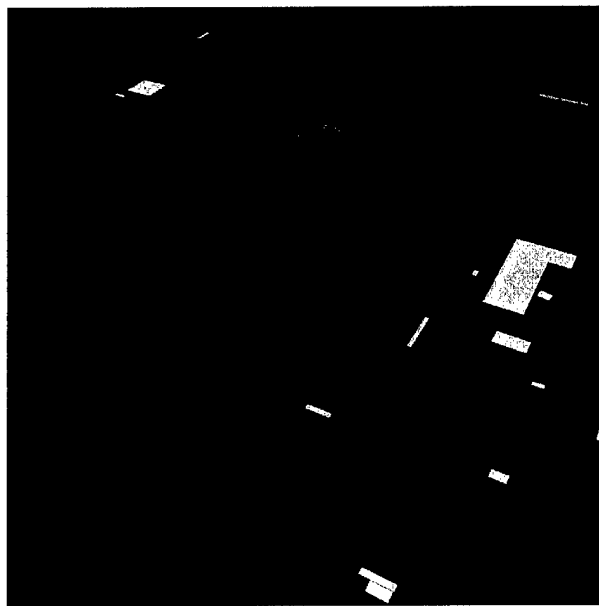


Figure 19: Surface material assignment after translation.

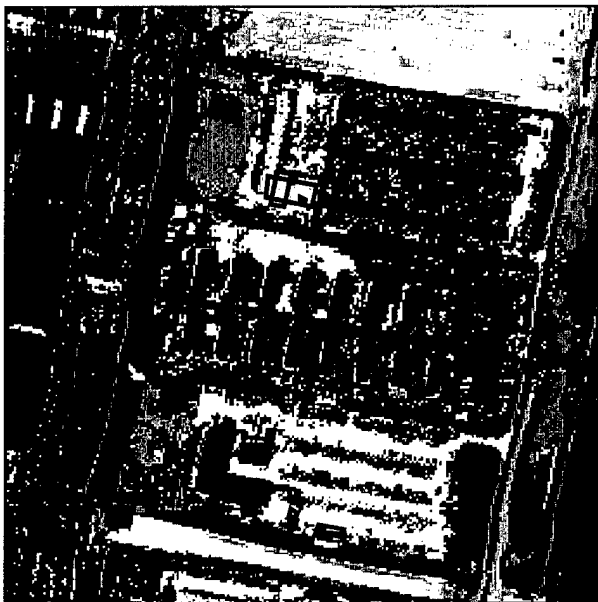


Figure 20: Refined classmap using discriminant analysis.

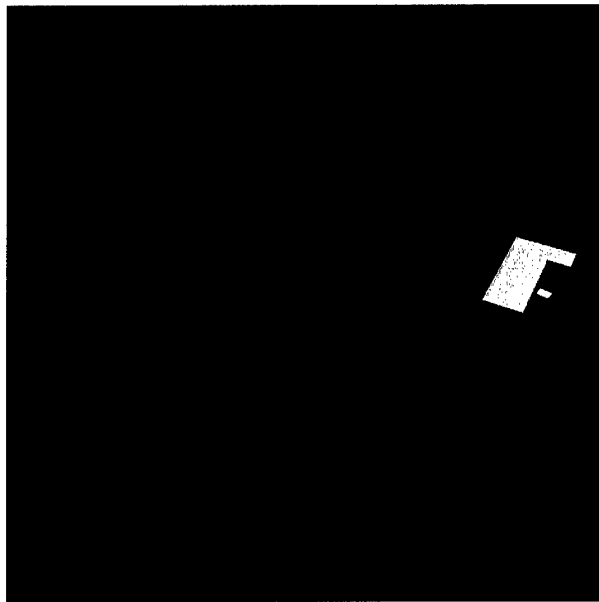
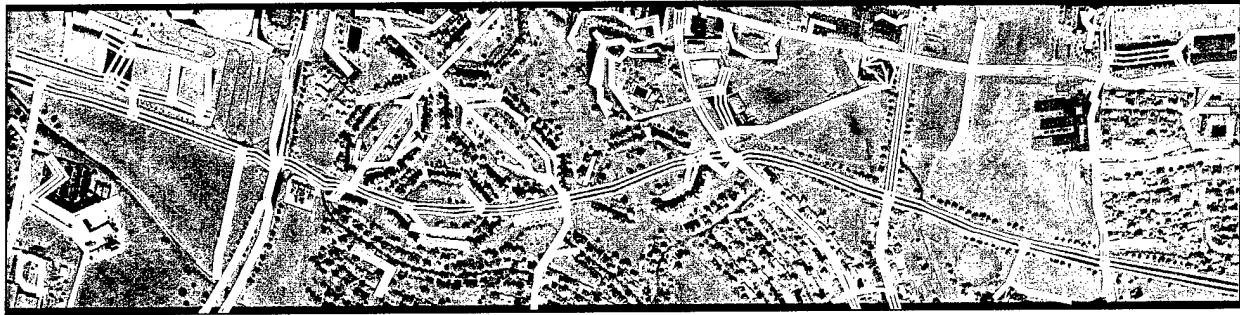
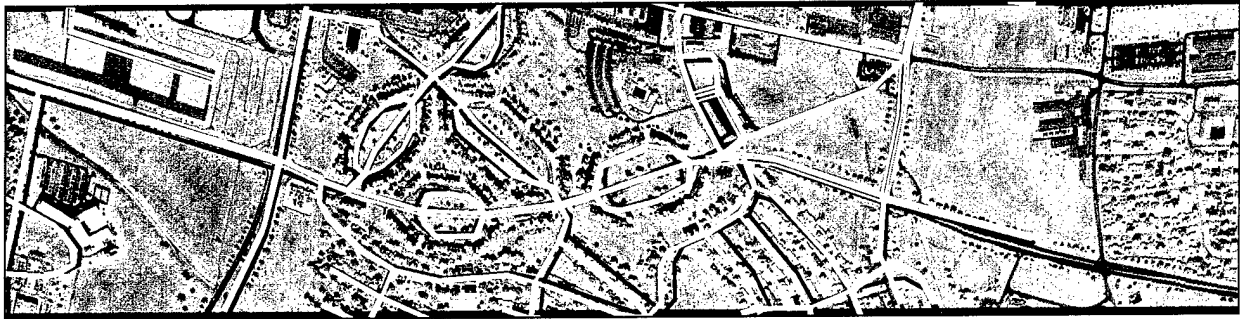


Figure 21: Final results for surface material assignment.



(a) Automatically generated road network extracted using HYDICE derived surface material classification.



(b) Reprojected DLG road network (blue) and automatically adjusted DLG road network with width and surface material attributes derived from HYDICE surface material classification (yellow).

Figure 22: Road network extracted using the surface material classification generated from the HYDICE data.

comprising the DLG road network. Once translated, we overlay the road network on the surface material map and compute a width for each of the road points. Surface material attributions also can be generated at this time. We can now generate full road models using the translated DLG centerlines and the computed width information. The registration refinement and width attribution steps are repeated again to ensure that the placement of the network is accurate given the improved width assignments.

Figure 22(b) shows a road network that has been adjusted using the process previously described. The new network seems to be better positioned on the image, and the computed width attributions appear to be close. Most problems are mismatches that can occur where there is more than one compatible surface material region that can be matched to the road, or when the match window is too small to allow the correct match to occur.

9. Conclusions

Our work under the APGD program has shown the applicability and suitability of hyperspectral data for surface material classification as input for visual simulation databases and land cover studies. In particular, it has been shown to be especially effective as a component for fusion-based cartographic feature extraction, in conjunction with stereo elevation or building and road extraction systems.